

LISTING OF THE CLAIMS

The following listing, if entered, replaces all prior versions of the claims in the present application.

1. **(Currently Amended)** A method comprising:
 identifying a plurality of secondary nodes, ~~wherein;~~
~~the identifying comprises~~ sending an update from a primary node to the plurality
 of secondary nodes, ~~and wherein~~
the update identifies at least one secondary node of the plurality of
secondary nodes to which the update will be sent;
in response to receiving the update from the primary node, causing the at
 least one secondary node ~~of the plurality of secondary nodes inserts to~~
insert the update in a respective log of updates, ~~and wherein~~
 each of the respective log of updates corresponds to a respective copy of
 the data, ~~and~~
the inserting the update in the respective log of updates is performed
at the secondary node;
in response to inserting the update in the respective log of updates, causing
the update to be copied from the respective log of updates to a storage
area at the least one secondary node;
sending an acknowledgement from the at least one secondary node to the
primary node, wherein
the acknowledgement indicates that the update has been received at
the least one secondary node;
 determining that all of the plurality of secondary nodes have acknowledged the
 update; and
in response to the determining, causing each secondary node of the at least one
 secondary node to clear the update from the respective log of updates by
 sending a notification to each of the plurality of secondary nodes once all
 of the plurality of secondary nodes have acknowledged the update,
 wherein
 [[said]] the clearing is performed in response to receiving the notification.

2. (Cancelled)
3. **(Currently Amended)** The method of claim 1 wherein **the** clearing the update from the respective log comprises updating a start-of-log pointer in the respective log.
4. (Previously Presented) The method of claim 1 wherein the clearing the update from the respective log comprises updating a pointer to a location in the respective log, wherein the pointer points to the location if the location contains a next update to clear.
5. (Previously Presented) The method of claim 1 further comprising:
determining that a location of a next update in a first respective log of updates to a first respective copy of the data at a first secondary node of the secondary nodes differs from a corresponding location of the next update in a second respective log of updates to a second respective copy of the data at a second secondary node of the secondary nodes; and
identifying a set of updates in the first respective log, wherein each update of the set of updates is not in the second respective log; and
synchronizing the first respective copy and the second respective copy by applying the set of updates to the second respective copy.
6. (Previously Presented) The method of claim 5 wherein the determining occurs when a primary node maintaining the data fails.
7. (Previously Presented) The method of claim 1 further comprising:
setting a sent indicator for the update for one of the plurality of secondary nodes when the update is sent to the one secondary node.
8. (Previously Presented) The method of claim 7 further comprising:
setting a received indicator for the update for the one secondary node when an acknowledgement of the update is received from the one secondary node.

9. (Previously Presented) The method of claim 8 wherein the sending the notification to each of the plurality of secondary nodes comprises determining that a respective sent indicator and a respective received indicator for the update are set for each of the plurality of secondary nodes.

10.-17. (Cancelled)

18. (Currently Amended) A computer-readable storage medium having a plurality of instructions embodied therein, wherein ~~[[said]]~~ the plurality of instructions are executable for:

identifying instructions to identify a plurality of secondary nodes, ~~wherein;~~
~~said identifying comprises~~ sending an update from a primary node to ~~[[said]]~~ the plurality of secondary nodes, ~~and wherein~~
the update identifies at least one secondary node of the plurality of secondary nodes to which the update will be sent;
in response to receiving the update from the primary node, causing the at
least one secondary node ~~of the plurality of secondary nodes inserts to~~
insert the update in a respective log of updates, ~~and wherein~~
each of the respective log of updates corresponds to a respective copy of the data, ~~and~~
the inserting the update in the respective log of updates is performed at the secondary node;
in response to inserting the update in the respective log of updates, causing the update to be copied from the respective log of updates to a storage area at the least one secondary node;
sending an acknowledgement from the at least one secondary node to the primary node, wherein
the acknowledgement indicates that the update has been received at the least one secondary node;
determining that all of the plurality of secondary nodes have acknowledged the update; and
in response to the determining, causing each secondary node of the at least one

secondary node to clear the update from the respective log of updates by sending instructions to send a notification to each of the plurality of secondary nodes once all of the plurality of secondary nodes have acknowledged the update, wherein
 [[said]] the clearing is performed in response to receiving the notification.

19. (Cancelled)
20. (Previously Presented) The computer-readable storage medium of claim 18 wherein the clearing instructions further comprise
 updating instructions to update a start-of-log pointer in the respective log.
21. (Previously Presented) The computer-readable storage medium of claim 18 wherein the clearing instructions further comprise
 updating instructions to update a pointer to a location in the respective log,
 wherein
 the pointer points to the location if the location contains a next update to clear.
22. (Previously Presented) The computer-readable storage medium of claim 18 further comprising:
 determining instructions to determine that a location of a next update in a first
 respective log of updates to a first respective copy of the data at a first
 secondary node of the secondary nodes differs from a corresponding
 location of the next update in a second respective log of updates to a
 second respective copy of the data at a second secondary node of the
 secondary nodes; and
 second identifying instructions to identify a set of updates in the first respective
 log, wherein
 each update of the set of updates is not in the second respective log; and
 synchronizing instructions to synchronize the first respective copy and the second
 respective copy by applying the set of updates to the second respective
 copy.

23. (Currently Amended) A computer system comprising:
- a processor for executing instructions, and
 - a memory to store the instructions, wherein the instructions comprise
 - identifying instructions to identify a plurality of secondary nodes to which an update to data is sent from a primary node, wherein the update identifies at least one secondary node of the plurality of secondary nodes to which the update will be sent;
 - in response to receiving the update from the primary node, causing the at least one secondary node ~~of the plurality of secondary nodes~~ to insert[[s]] the update in a respective log of updates to a respective copy of the data, wherein
 - each of the respective log of updates corresponds to a respective copy of the data, and
 - the inserting the update in the respective log of updates is performed at the secondary node;
 - in response to inserting the update in the respective log of updates, causing the update to be copied from the respective log of updates to a storage area at the least one secondary node;
 - sending an acknowledgement from the at least one secondary node to the primary node, wherein
 - the acknowledgement indicates that the update has been received at the least one secondary node;
 - determining that all of the plurality of secondary nodes have acknowledged the update; and
 - in response to the determining, causing each secondary node of the at least one secondary node to clear the update from the respective log of updates by sending instructions to send a notification to each of the plurality of secondary nodes when all of the plurality of secondary nodes have acknowledged the update, wherein [[said]] the clearing is performed in response to receiving the notification.

24. (Cancelled)
25. (Previously Presented) The computer system of claim 23 wherein the instructions further comprise
- determining instructions to determine that a location of a next update in a first respective log of updates to a first respective copy of the data at a first secondary node of the secondary nodes differs from a corresponding location of the next update in a second respective log of updates to a second respective copy of the data at a second secondary node of the secondary nodes; and
- second identifying instructions to identify a set of updates in the first respective log, wherein
- each update of the set of updates is not in the second respective log; and
- synchronizing instructions to synchronize the first respective copy and the second respective copy by applying the set of updates to the second respective copy.
26. (Previously Presented) The method of claim 1, further comprising:
- in response to the identifying, incrementing a regional counter stored on the primary node by a number of secondary nodes to which the update is sent, wherein
- the regional counter is a number of secondary nodes from which an acknowledgement to the update is to be received;
- in response to receiving an acknowledgement from a secondary node among the plurality of secondary nodes to which the update is sent, decrementing the regional counter; and
- in response to the regional counter reaching a value prior to the incrementing, determining that each of the plurality of secondary nodes has acknowledged the update.

27. (Previously Presented) The computer-readable storage medium of claim 18, further comprising:

incrementing instructions to increment a regional counter stored on the primary node by a number of secondary nodes to which the update is sent, in response to the identifying, wherein
the regional counter is a number of secondary nodes from which an acknowledgement to the update has not been received;
decrementing instructions to decrement the regional counter, in response to receiving an acknowledgement from a secondary node among the plurality of secondary nodes to which the update is sent; and
determining instructions to determine that each of the plurality of secondary nodes has acknowledged the update.

28. (Previously Presented) The computer system of claim 23, wherein the instructions further comprise:

incrementing instructions to increment a regional counter stored on the primary node by a number of secondary nodes to which the update is sent, in response to the identifying, wherein
the regional counter is a number of secondary nodes from which an acknowledgement to the update has not been received;
decrementing instructions to decrement the regional counter, in response to receiving an acknowledgement from a secondary node among the plurality of secondary nodes to which the update is sent; and
determining instructions to determine that each of the plurality of secondary nodes has acknowledged the update.